Deloitte.

Reprinted by permission of Harvard Business Review Press. Excerpted from Keeping Up with the Quants by Thomas H. Davenport and Jinho Kim.

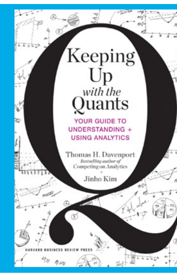

Types of analytics stories

There are at least six types of quantitative analytical stories; each of them is described below, along with an example or two.

The CSI story

Some quantitative analyses are like police procedural television programs; they attempt to solve a business problem with quantitative analysis. Some operational problem crops up, and data are used to confirm the nature of the issue and find the solution. This situation often does not require deep statistical analysis, just good data and reporting approaches. It is often encountered in online businesses, where customer clickstreams provide plenty of data—often too much—for analysis.

One expert practitioner of the CSI story approach is Joe Megibow, vice president and general manager of online travel company Expedia's U.S. business. Joe was previously a Web analytics maven—and he still is—but his data-based problem-solving approaches have led to a variety of impressive promotions.

Many of the Expedia investigations involve understanding the reasons behind lost online sales. One particular CSI story involved lost revenue on hotel payment transactions. Analysis of data suggested that after a customer had selected a hotel, filled in the travel and billing information, then clicked the "Buy Now" button, a percentage of the sales transactions were not being completed successfully. Megibow's team investigated the reason for the failures, again using Web metrics data and server log files throughout the process.

Apparently, the "Company" field under the customer's name was causing a problem. Some customers interpreted it as the name of the bank that supplied their credit card, and then they also supplied the bank's address in the billing address fields. This caused the transaction to fail with the credit card processor. Simply removing the "Company" field immediately raised profits for Expedia by $12 million. Megibow says that Expedia has explored many of these CSI-like stories, and they almost always yield substantial financial or operational benefits.

Sometimes the CSI stories do involve deeper quantitative and statistical analysis. One member of Megibow's team was investigating which customer touchpoints were driving online sales transactions. The analyst used the Cox regression model—an approach originally used to determine which patients would die and which would live over certain time periods—of "survival analysis." The analysis discovered that the simpler prior models were not at all correct about what marketing approaches were really leading to a sale. Megibow commented, "We didn't know we were leaving money on the table."[1]

1 Study on data scientists: Thomas H. Davenport, The Human Side of Big Data and High-Performance Analytics (sponsored by SAS and EMC), http://www.sas.com/ reg/gen/corp/2154478; study on big data in big companies: Thomas H. Davenport and Jill Dyché, Big Data in Big Companies Research Report (sponsored by SAS), www.sas.com/reg/gen/corp/2266746; study of big data in the travel industry: Thomas H. Davenport, At the Big Data Crossroads: Turning Toward a Smarter Travel Experience (sponsored by Amadeus), www.amadeus.com/bigdata; study of data discovery: sponsored by Teradata Aster (not yet completed at the time of this book's publication).

The Eureka story

The Eureka story is similar to the CSI story, except that it typically involves a purposeful approach to a particular problem (as opposed to stumbling over the problem) to examine a major change in an organization's strategy or business model. It tends to be a longer story with a greater degree of analysis over time. Sometimes Eureka stories also involve other analytical story types, just because the results are so important to the organizations pursuing them.

At Expedia again, for example, one Eureka story involved eliminating change/cancel fees from online hotel, cruise, and car rental reservations. Until 2009, Expedia and its competitors all charged up to $30 for a change or cancellation—above and beyond the penalties the hotel imposed. Expedia and other online bookers' rates were typically much lower than booking directly with a hotel, and customers were willing to tolerate change/cancel fees.

However, by 2009 it had become apparent that the fees had become a liability. Expedia's rates were closer to those of the hotels' own rates, so the primary appeal of Expedia had become convenience— and change/cancel fees were not convenient. Analysts looked at customer satisfaction rates, and they were particularly low for customers who had to pay the fees. Expedia's call center representatives were authorized to waive the change/cancel fees for only one reason: a death in the customer's family. A look at the number of waivers showed double-digit growth for the past three years. Either there was a death epidemic, or customers had figured out they could get their money back this way.

Expedia executives realized the market had changed, but change/ cancel fees represented a substantial source of revenue. They wondered if the fees were eliminated, would conversion (completed sale) rates go up? In April of 2009, they announced a temporary waiver of fees for the month (a bit of a mad scientist testing story, described below). Conversion rates immediately rose substantially. Executives felt that they had enough evidence to discontinue the fees, and the rest of the industry followed suit.

Across town in Seattle lies Zillow, a company that distributes information about residential real estate. Zillow is perhaps best known to quant jocks for its "Zestimates," a proprietary algorithm that generates estimates of home values. But, like Expedia, Zillow's entire culture is based on data and analysis—not surprisingly, since the company was founded by Rich Barton, who also founded Expedia.

One of Zillow's Eureka stories involved a big decision to change how it made its money from relationships with real estate agents. Zillow began to work with agents in 2008, having previously been focused on consumers. One aspect of its agent-related business model was selling advertising by agents and delivering leads to them. Zillow charged the agents for the leads, but the value per lead was not enough in the view of executives. Chloe Harford, a Zillow executive who heads product management and strategy, was particularly focused on figuring out the right model for increasing lead value and optimizing the pricing of leads.

Harford, who has a PhD in volcanology, or the study of volcanoes, is capable of some pretty sophisticated mathematical analysis. However, she and her colleagues initially relied on what she calls "napkin math" to explore other ways to generate more leads and price them fairly to agents. In April 2010, Zillow created a new feature—immediately copied by competitors—involving selling advertising to agents. It created many more customer contacts than before, and allowed the consumer to contact the agent directly. Zillow also introduced a sophisticated algorithm for pricing leads to agents that attempts to calculate the economic value of the lead, with an estimate of conversion rates. Competitors also do this to some degree, but probably not to the level of sophistication that Zillow does. The leads and pricing of them are so important that Harford and her colleagues frequently test different approaches of them with some of the Mad Scientist testing approaches described below. In short, Zillow's Eureka stories are intimately tied into its business model and its business success.

The mad scientist story

We're all familiar with the use of scientific testing in science-based industries such as pharmaceuticals. Drug companies test their products on a group of test subjects, while giving a placebo to members of a control group. They pay careful attention to ensure that people are randomly assigned to either the test or control group, so there are no major differences between the groups that might impact the drug's effectiveness. It's a powerful analytical tool because it's usually as close as we can come to causation—the knowledge that what is being tested in the test group is driving the outcome in a causal fashion.

Rigorous testing is no longer just the province of white-coated scientists; it is now an analytical approach that every large organization can employ. There is broadly available software that leads managers or analysts through the testing process. Companies can now base important decisions on real, scientifically valid experiments. In the past, any foray into randomized testing (the random assignment to groups that we mentioned above) meant employing or engaging a PhD in statistics or a "design of experiments" expert. Now, a quantitatively trained MBA can oversee the process, assisted by software that will help determine what sizes of groups are necessary, which sites to use for testing and controls, and whether any changes resulting from experiments are statistically significant.

The mad scientist stories are particularly well suited to organizations like retailers (that have a lot of stores) and banks (that have a lot of branches). That makes it easy to try things out in some locations and use others as controls. It's also quite easy to do testing on websites, where you can send some customers to one version of a Web page, send other customers to a different version, and see if the results are significantly different (called A/B testing in the Web analytics field).

Some examples of mad scientist stories include:[2]

- Do lobster tanks sell more lobsters at Food Lion supermarkets? The answer is apparently yes if the store was one in which customers already bought lobsters (i.e., they were relatively upscale), and no if the store didn't attract lobster-buying customers to begin with.
- Does a Sears store inside a Kmart sell more than all-Kmart? Sears Holdings chairman Eddie Lampert is a big fan of randomized testing and has tested a variety of such combinations. We don't know the answer to this particular question, but we're guessing that if the answer were a definitive yes, we would have seen a lot more of these blended stores.
- Are the best sales results at the Red Lobster seafood restaurant chain achieved from a low-, medium-, or high-cost remodel of restaurants-and should the exterior or the interior be the primary focus? The result, according to Red Lobster executives, was that the medium-cost interior remodel paid off best. Exterior remodels brought a lot of new customers in, but if they saw that the interiors hadn't been redone as well, they didn't come back.

2 NewVantage Partners, Big Data Executive Survey: Themes and Trends, 2012, http://newvantage.com/data-management/. I am an adviser to NewVantage Partners.

The survey story

Surveys are a classic method of quantitative research. The survey analyst observes phenomena that have already happened or are happening now. The analyst doesn't try to manipulate the outcome—only to observe, codify, and analyze it. Typically the surveyor seeks to understand what traits or variables observed in the survey are statistically related to other traits. The simplest example would be if we asked a sample of customers of a particular product various things about themselves, including demographic information like gender and age. If we also asked what products they liked, we could then determine whether men like certain products more than women, or whether certain products are more likely to be liked by younger people.

Surveys are popular and relatively easy to carry out. However, we have to remember that the results and stories based on them can vary considerably based on how questions are asked and how they vary (or not) over time. For example, the U.S. Census has worked for literally decades on questions about the race of U.S. citizens. The number of racial categories in census surveys keeps expanding; in the 2010 census there were fifteen choices, including "some other race." That was a popular choice for the more than 50 million Latino U.S. citizens, 18 million of whom checked the "Other" box.[3] If there is that much confusion about race, imagine what difficulties survey researchers can have with slippery topics such as politics, religion, social attitudes, and sexual behavior.

We also have to remember that just because two variables in a survey analysis are related, they may not be causally related. There may well be other variables that you're not looking at that might be the causal factor driving the phenomena you care about.

Survey stories often involve asking people about their beliefs and attitudes, but they don't have to involve people. Take, for example, this survey of airplanes conducted during World War II, related in a classic statistics textbook:

During the Second World War it was necessary to keep planes in action as much as possible, so it was decided to see if the number of time-consuming engine overhauls could be reduced without risk. A retrospective survey was made of planes that were lost, and contrary to all expectations, it was found that the number of planes lost as a result of engine troubles was greatest right after overhaul, and actually decreased as the time since overhaul grew longer. This result led to a considerable increase in the intervals between overhauls, and needless to say, to important revisions in the manner of overhauling to make sure that all those nuts and bolts were really tightened up properly.[4]

If you're planning to do or analyze a survey, make sure that you've thought very carefully about the meanings of your survey questions or variables. A variable is any measured characteristic, with two or more levels or values, of properties of people, situations, and behaviors. Gender, test scores, room temperature, love, happiness, and team cohesiveness are good examples of variables.

Also, it's important to ensure that your survey sample is representative of the population you want to study. How you perform the survey can affect the sample. For example, if you want to survey young people's attitudes or behaviors, don't hire a survey firm that only contacts the members of the sample through landline telephones. That's a very typical approach, but we all know that many young people don't have, and don't ever intend to have, a landline. So they would be underrepresented in a sample that employs only landlines.[5]

---

3 Dan Power's A Brief History of Decision Support Systems has more detail on some of the early terminology; see http://dssresources.com/history/dsshistory.html.

4 The 2.5-quintillion-byte estimate comes from IBM, "What Is Big Data? Bringing Big Data to the Enterprise," www.ibm.com.

5 H. James Wilson, "You, by the Numbers," Harvard Business Review, September 2012, 119–122.

The prediction story

Prediction stories are all about anticipating what will happen in the future. While it's pretty difficult to get good data about the future, taking data about the past and understanding the factors that drive past events is pretty straightforward for quantitative analytics. Typically this is referred to as predictive analytics or predictive modelling.

There are a variety of prediction stories that an analyst can construct. Below is a sample of possibilities; note how specific they are:

- Offer response: Which customers will respond to an e-mail of a free shipping offer within two business days with a purchase of $50 or more?
- Cross-sell/upsell: Which checking account customers with account balances over $2,000 will purchase a one-year CD with an interest rate of 1.5 percent, responding within one month, given a mail solicitation?
- Employee attrition: Which employees of more than six months who haven't yet signed up for the 401(k) program will resign from their jobs within the next three months?

There are many other predictive analytics possibilities. In business, a common approach to prediction is to determine what offer the customer is most likely to accept. The most sophisticated versions of this "next best offer" analytics are increasingly automated; no human needs to see the offer before it is made available to the customer, and there can be hundreds or thousands of different offers.

Microsoft, for example, has an incredible ability to dynamically tailor "offers" for its Bing search engine (the product is free, so Microsoft is just trying to get you to use it). The offers tempt you to try out Bing, to create a Bing search bar on your browser, to try a particular Bing feature, and so forth. The customization of the offer is based on a variety of factors—including your location, age, gender, and recent online activity—that it can determine from your cookies and other sources. If you have signed up for Microsoft Passport, the company has even more information about you that allows for targeting the offers even more effectively. Microsoft is able (facilitated by the Infor Epiphany Interaction Advisor software they use) to instantly compose a targeted e-mail the moment you click on an offer in your inbox; it all takes about 200 milliseconds. Microsoft says it works extremely well to lift conversion rates.

Often, prediction stories can be a bit of a fishing expedition. We don't know exactly what factors will allow us to predict something, so we try a lot of them and see what works. Sometimes the results are unexpected. For example, in the Microsoft Bing offers we've just described, the number of Microsoft Messenger buddies you have turns out to be a good predictor of whether you'll try out Bing.

At Google, the company wanted to predict what employee traits predicted high performance. Some analysis determined that the factors Google was originally using—grades in college and interview ratings—were poor predictors of performance. Since they weren't sure what factors would be important, they asked employees to answer a three-hundred-question survey. As Laszlo Bock, the head of People Operations at Google, noted: "We wanted to cast a very wide net. It is not unusual to walk the halls here and bump into dogs. Maybe people who own dogs have some personality trait that is useful."[6]

Bringing pets to work didn't prove to predict much of anything, but Google did find some unexpected predictors. For example, whether a job applicant had set a world or national record or had started a nonprofit organization or club were both associated with high performance. Google now asks questions about experiences like these on its online job interviews.

Of course, if the factors that predict something make no sense at all, it's a good idea to go back and recheck your data and your analysis. But actually looking at some data can outperform a human futurist's predictions much of the time. As a caution, remember that predictive stories use data from the past to tell stories about the future. If something in the world has changed since you did your analysis, the predictions may no longer hold.

6 Stephen Wolfram, "The Personal Analytics of My Life," blog post, March

The “here’s what happened” story

Stories that simply tell what happened using data are perhaps the most common of all. They provide the facts—how many products were sold when and where, what were the financials that were achieved last quarter, how many people did we hire last year. Since they are reporting-oriented stories that often don’t use sophisticated math, it might seem that they would be easy to tell. However, the great rise in data within today’s organizations has been mirrored by a similar rise in reports based on data. Therefore, it’s sometimes difficult to get the attention of the intended audience for the reports you create or distribute.

This type of story is particularly well suited to visual displays of information. Suffice it to say that if you are providing reports in rows and columns of numbers, you aren’t likely to get the attention you need. Many of us even tire today of colorful graphs and charts, but most people would say they are more worthy of attention than numbers on a page.